arXiv:1702.05327v1 [cs.LG] 17 Feb 2017

Anchored Regression: Solving Random Convex Equations via Convex Programming

Sohail Bahmani
sohail.bahmani@ece.gatech.edu

Justin Romberg
jrom@ece.gatech.edu

February 20, 2017

Abstract

We consider the problem of estimating a solution to (random) systems of equations that involve convex nonlinearities which has applications in machine learning and signal processing. Conventional estimators based on empirical risk minimization generally lead to non-convex programs that are often computationally intractable. We propose *anchored regression*, a new approach that utilizes an *anchor vector* to formulate an estimator based on a simple convex program. We analyze accuracy of this method and specify the required sample complexity. The proposed convex program is formulated in the natural space of the problem rather than a *lifted* domain, which makes it computationally favorable. This feature of anchored regression also provides great flexibility as structural priors (e.g., sparsity) can be seamlessly incorporated through convex regularization. We also provide recipes for constructing the anchor vector from the data.

1 Introduction

We consider the problem of (approximately) solving a system of nonlinear equations with convex nonlinearities. In particular, we observe

$$\begin{aligned} y_1 &= f_1(\boldsymbol{x}_\star) + \xi_1 \\ y_2 &= f_2(\boldsymbol{x}_\star) + \xi_2 \\ &\vdots \\ y_M &= f_M(\boldsymbol{x}_\star) + \xi_M \end{aligned}, \tag{1}$$

where $\boldsymbol{x}_\star \in \mathbb{R}^N$ is the ground truth, the known functions $f_1,\ f_2,\ \ldots,\ f_M$ are twice-differentiable and convex, and $\xi_1,\ \xi_2,\ \ldots,\ \xi_M$ are additive noise terms. Given these observations, we formulate an estimator for $\boldsymbol{x}_\star$ via a convex program that finds candidate solutions that are consistent with the observation and closely resemble any structure of the ground truth that is induced through a convex regularizaion term. Assuming that we have the computational oracles of f_ms and their derivatives required for convex optimization, the proposed estimator

can be computed efficiently. Our main results give error bounds on the quality of the produced estimate, specified by the standard Euclidean distance to $\boldsymbol{x}_\star$, under a model where the f_ms are drawn at random from a set $\mathcal{F}$ according to some probability law. The estimation error, and the sufficient number of equations M needed to achieve it, depend on a *Rademacher complexity* that involves the probability law imposed on $\mathcal{F}$ through the gradients $\nabla f_m(\boldsymbol{x}_\star)$ as well as the geometry induced by the regularization.

To illustrate broad applications of this observation model, we provide two motivating examples. The first example is learning with a *single hidden-layer neural network* [see, e.g., 7, 21]. Given an activation function $\sigma : \mathbb{R} \to \mathbb{R}$ and prescribed weights $w_1, w_2, \ldots, w_K \geq 0$ for the hidden layer, the output of a single hidden-layer neural network can be expressed as $f_m(\boldsymbol{x}) = \sum_{k=1}^K w_k \sigma(\boldsymbol{a}_m^\mathsf{T} \boldsymbol{x}_k)$, where $\boldsymbol{x} = \begin{bmatrix} \boldsymbol{x}_1^\mathsf{T} & \boldsymbol{x}_2^\mathsf{T} & \cdots & \boldsymbol{x}_K^\mathsf{T} \end{bmatrix}^\mathsf{T}$ denotes the input layer's weight parameters and $\boldsymbol{a}_m$ denotes the mth data sample. For convex activation functions such as $\sigma(u) = \log_2(1 + e^u)$, the functions $f_m(\boldsymbol{x})$ are also convex and the model assumed in (1) applies. The second example is the problem of *phase retrieval* from signal processing and computational imaging. The goal in phase retrieval is to estimate an signal/image (up to a global phase) from intensity of some linear observations. More precisely, given (random) *measurement vectors* $\boldsymbol{a}_1, \boldsymbol{a}_2, \ldots, \boldsymbol{a}_M \in \mathbb{C}^N$ the goals is to estimate $\boldsymbol{x}_\star \in \mathbb{C}^N$ from noisy intensity measurements $y_m = |\boldsymbol{a}_m^* \boldsymbol{x}_\star|^2 + \xi_m$ for $m = 1, 2, \ldots, M$. Clearly, this observation model is a special case of (1) with $f_m(\boldsymbol{x}) = |\boldsymbol{a}_m^* \boldsymbol{x}|^2$ being the randomly drawn functions.[1]

In this paper, we propose *anchored regression* to estimate $\boldsymbol{x}_\star$ in the parametric model described by (1). An *anchor vector* $\boldsymbol{a}_0 \in \mathbb{R}^N$ is a unit vector (i.e., $\|\boldsymbol{a}_0\|_2 = 1$) that obeys

$$\langle \boldsymbol{a}_0, \boldsymbol{x}_\star \rangle \geq \delta \|\boldsymbol{x}_\star\|_2 , \tag{2}$$

for an absolute constant $\delta \in (0, 1]$. In words, the anchor vector has a non-vanishing correlation with a ground truth. Given an anchor vector $\boldsymbol{a}_0$, our proposed estimator for (1) is the convex program

$$\begin{aligned} &\operatorname*{argmax}_{\boldsymbol{x}} \ \langle \boldsymbol{a}_0, \boldsymbol{x} \rangle - \Omega(\boldsymbol{x}) \\ &\text{subject to } R_M^+(\boldsymbol{x}) \leq R_M^+(\boldsymbol{x}_\star) + \varepsilon , \end{aligned} \tag{3}$$

where $\Omega(\boldsymbol{x})$ is a convex regularizer, $\varepsilon > 0$ is a small constant, and $R_M^+(\cdot)$ is the empirical one-sided error defined as

$$R_M^+(\boldsymbol{x}) \stackrel{\text{def}}{=} \frac{1}{M} \sum_{m=1}^M (f_m(\boldsymbol{x}) - y_m)_+ ,$$

which is also convex. The function $(\cdot)_+$ denotes the positive part. Of course, the value of $R_M^+(\boldsymbol{x}_\star)$ is unknown in general, but we assume it can be estimated and ε in (3) represents the error in the estimate.

1.1 How does the anchor help?

Because the functions $f_1, f_2, \ldots, f_M$ are convex, in general, the system of equations (1) may have multiple "equivalent" solutions. For convex empirical risk functions, having a

[1] While treatment of complex-valued vectors is slightly different from the real-valued vectors, this difference does not fundamentally change the formulation and the proposed estimator.

non-unique ground truth generally leads to degenerate solutions in the convex hull of the equivalent ground truth parameters. To avoid this problem, in some cases we can *lift* the parameters to a higher-dimensional space where the equivalent solutions collapse to a single solution. However, this approach has limited scope of applicability and usually leads to a significant increase in the computational cost.

The use of an anchor vector is a versatile method to resolve the ambiguity caused by a non-unique ground truth without incurring a significant computational cost. To illustrate this point, consider the phase retrieval problem discussed above. Suppose that $f_m(\boldsymbol{x}) = |\boldsymbol{a}_m^* \boldsymbol{x}|^2$ for randomly drawn vectors $\boldsymbol{a}_m$. There is an inherent solution ambiguity in the phase retrieval problem due to the fact that a global change of phase does not affect the observations; for any value of ϕ the vectors $e^{\imath\phi}\boldsymbol{x}_\star$ produce the same observations as $\boldsymbol{x}_\star$. It suffices to find only one of the equivalent solutions. Prior convex relaxations for the phase retrieval problem use a lifting transform, such as $\boldsymbol{x}\boldsymbol{x}^* \mapsto \boldsymbol{X}$, to embed the unknowns in a higher dimensional space of positive semidefinite matrices. This transform effectively converts the equations that are quadratic in the natural variable into equations that are linear in the lifted variable. Then, the estimator of the lifted solution can be formulated as a semidefinite program that approximates a rank-one solution to the lifted linear equations through some variant of trace-norm regularization. This technique is proved to achieve optimal sample complexity under different measurement distributions, but it is not practically scalable due the prohibitive computational cost of semidefinte programming.

In Section 3 we describe some schemes and conditions by which anchor vectors can be constructed from the existing random observations.

1.2 Related work

The proposed observation model (1) applies to regression with linear *predictors* where we have $f_m(\boldsymbol{x}) = \mu(\boldsymbol{a}_m^{\mathsf{T}} \boldsymbol{x})$ for the (random) data sample $\boldsymbol{a}_m$ and a known convex function μ. For example, our framework is an alternative to maximum likelihood estimation in many instances of *generalized linear models* (GLM) [14] with μ being the *mean function* (or the *inverse link function*) of the GLM. In such cases, the mean function is either convex or it has a property (e.g., it is concave, log-convex, or log-concave) that allows us to easily convert it to a convex function. Another relevant model is the semiparametric *single index model* [see, e.g., 8] where, again, the model assumes linear predictors, but the nonlinear function μ is not known. Under some regularity assumption on μ, in [19, 20] simple estimators based on convex programming are shown to produce accurate estimates (up to some scaling factor). Because μ is unknown, our framework does not apply to the single index model. However, there are interesting instances of (1) (e.g., phase retrieval) where the results of [19, 20] do not apply as the assumed regularity conditions lead to trivial bounds.

The most relevant results to our work are the recent methods proposed for phase retrieval in [1] and independently in [3] that exploit anchor vectors. As mentioned above, the phase retrieval problem can be described by the model (1) with $f_m(\boldsymbol{x}) = |\boldsymbol{a}_m^* \boldsymbol{x}|^2$ for some random measurement vector $\boldsymbol{a}_m$. It is shown in [1, 3] that, using an anchor, the ground truth $\boldsymbol{x}_\star$ can be estimated accurately and with optimal sample complexity through a convex program analogous to (3). The analyses in [3] and [1] differ in that the former assumes the anchor is independent of the measurements whereas the latter does not make this assump-

tion. Alternative proofs and variations of this phase retrieval method also appeared later in [6, 4, 5].

2 Main Result

Let $\partial\Omega(\boldsymbol{x}_\star)$ denote the *subdifferential* of $\Omega(\cdot)$ at $\boldsymbol{x}_\star$. To show that (3) produces an accurate estimate of $\boldsymbol{x}_\star$, it suffices to show that the set of ascent directions,

$$\mathcal{A} = \{\boldsymbol{h} \,:\, \langle \boldsymbol{a}_0 - \boldsymbol{g}, \boldsymbol{h}\rangle \geq 0, \text{ for all } \boldsymbol{g} \in \partial\Omega(\boldsymbol{x}_\star)\}\,,$$

does not contain any vector with large ℓ_2 norm that is consistent with the constraint. As explained below in Section 3, the anchor $\boldsymbol{a}_0$ may be constructed from the existing observations. The dependence of the set $\mathcal{A}$ on the data through $\boldsymbol{a}_0$ makes the analysis difficult. To avoid this dependence, we require the sufficient condition to hold on some expansion of $\mathcal{A}$ that is decoupled from $\boldsymbol{a}_0$. Using the property of the anchor vector (2) we can write

$$\begin{aligned}
\langle \boldsymbol{a}_0, \boldsymbol{h}\rangle &= \langle \boldsymbol{a}_0 - \frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2}, \boldsymbol{h}\rangle + \langle \frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2}, \boldsymbol{h}\rangle \\
&\leq \left\| \boldsymbol{a}_0 - \frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2} \right\|_2 \|\boldsymbol{h}\|_2 + \langle \frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2}, \boldsymbol{h}\rangle \\
&\leq \sqrt{2(1-\delta)}\,\|\boldsymbol{h}\|_2 + \langle \frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2}, \boldsymbol{h}\rangle\,.
\end{aligned}$$

This implies the inclusion

$$\mathcal{A} \subset \mathcal{A}_\delta \overset{\text{def}}{=} \left\{ \boldsymbol{h} \,:\, \sqrt{2(1-\delta)}\,\|\boldsymbol{h}\|_2 + \langle \frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2} - \boldsymbol{g}, \boldsymbol{h}\rangle \geq 0, \text{ for all } \boldsymbol{g} \in \partial\Omega(\boldsymbol{x}_\star) \right\}. \tag{4}$$

It suffices to show that $\mathcal{A}_\delta$ does not contain any (relatively) large ℓ_2 ball. Our main theorem below provides a sample complexity for establishing the desired sufficient condition and thus accuracy of (3).

Theorem 1. *Let $\mathcal{A}_\delta$ be defined as (4). For a set $\mathcal{H} \subset \mathbb{R}^N$ define the Rademacher complexity*

$$\mathfrak{C}_M(\mathcal{H}) \overset{\text{def}}{=} \mathbb{E} \sup_{\boldsymbol{h}\in\mathcal{H}} \frac{1}{\sqrt{M}} \sum_{m=1}^{M} \epsilon_m \langle \nabla f_m(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\,, \tag{5}$$

where $\epsilon_1, \epsilon_2, \ldots, \epsilon_M$ are i.i.d. Rademacher random variables independent of everything else . Furthermore, for a positive constant τ, let

$$p_\tau(\mathcal{A}_\delta) = \inf_{\boldsymbol{h}\in\mathcal{A}_\delta} \mathbb{P}\left(\langle \nabla f(\boldsymbol{x}_\star), \boldsymbol{h}\rangle \geq \tau \|\boldsymbol{h}\|_2\right). \tag{6}$$

For any $t > 0$, if

$$M > \left(\frac{2\mathfrak{C}_M(\mathcal{A}_\delta) + t}{\tau p_\tau(\mathcal{A}_\delta)} \right)^2, \tag{7}$$

then with probability $\geq 1 - e^{-2Mt^2/\tau^2}$ any solution $\widehat{\boldsymbol{x}}$ of (3) obeys

$$\|\widehat{\boldsymbol{x}} - \boldsymbol{x}_\star\|_2 \leq \left(\frac{1}{M}\sum_{m=1}^{M} |\xi_m| + \varepsilon \right) \Big/ \left(\tau p_\tau(\mathcal{A}_\delta) - \frac{2\mathfrak{C}_M(\mathcal{A}_\delta) + t}{\sqrt{M}} \right).$$

The proof, provided below in Section 4, is based on the idea of *small-ball method* introduced in [10, 16] and further developed in [17, 11, 12]. Note that instead of (4) we could have used a slightly better approximation given by

$$\mathcal{A} \subseteq \left\{ \boldsymbol{h} \,:\, \sqrt{1-\delta^2}\,\|\boldsymbol{h}_\perp\|_2 + \langle \frac{\delta\,\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2} - \boldsymbol{g}, \boldsymbol{h}\rangle \geq 0, \text{ for all } \boldsymbol{g} \in \partial\Omega\left(\boldsymbol{x}_\star\right) \right\},$$

where $\boldsymbol{h}_\perp$ denotes the projection of $\boldsymbol{h}$ onto the hyperplane orthogonal $\boldsymbol{x}_\star$. While this approximation would improve the dependence of our result on δ, we prefer (4) merely for simpler notation and derivations.

Bounding the Rademacher Complexity $\mathfrak{C}_M\left(\mathcal{A}_\delta\right)$

To make the result of Theorem 1 more explicit we consider the special cases of unstructured regression (i.e., $\boldsymbol{x}_\star$ is arbitrary and $\Omega\left(\boldsymbol{x}\right) = 0$) and sparse regression (i.e., $\boldsymbol{x}_\star$ is sparse and $\Omega\left(\boldsymbol{x}\right) = \lambda\,\|\boldsymbol{x}\|_1$ for some $\lambda > 0$). The specific choice of $\Omega\left(\boldsymbol{x}\right)$ helps to simplify the Rademacher complexity $\mathfrak{C}_M\left(\mathcal{A}_\delta\right)$. Of course, $\mathfrak{C}_M\left(\mathcal{A}_\delta\right)$ would depend on the law of $f_1, f_2, \ldots, f_M$ as well.

Unstructured regression: In the first example, we approximate $\mathfrak{C}_M\left(\mathcal{A}_\delta\right)$ when no regularization is applied, i.e., $\Omega\left(\boldsymbol{x}\right) = 0$. In this case, (4) reduces to

$$\mathcal{A}_\delta = \left\{ \boldsymbol{h} \,:\, \sqrt{2\left(1-\delta\right)}\,\|\boldsymbol{h}\|_2 + \langle \frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2}, \boldsymbol{h}\rangle \geq 0 \right\},$$

which we approximate by the entire space (i.e., $\mathcal{A}_\delta \subseteq \mathbb{R}^N$). Because $\mathfrak{C}_M\left(\cdot\right)$ is monotonic, we obtain

$$\begin{aligned}
\mathfrak{C}_M\left(\mathcal{A}_\delta\right) &\leq \mathfrak{C}_M\left(\mathbb{R}^N\right) \\
&= \mathbb{E} \sup_{\boldsymbol{h}} \frac{1}{\sqrt{M}} \sum_{m=1}^{M} \epsilon_m \langle \nabla f_m\left(\boldsymbol{x}_\star\right), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2} \rangle \\
&= \mathbb{E} \left\| \frac{1}{\sqrt{M}} \sum_{m=1}^{M} \epsilon_m \nabla f_m\left(\boldsymbol{x}_\star\right) \right\|_2 .
\end{aligned}$$

Applying the Cauchy-Schwarz for the square-root function which is concave, we deduce that

$$\begin{aligned}
\mathfrak{C}_M\left(\mathcal{A}_\delta\right) &\leq \sqrt{\mathbb{E} \left\| \frac{1}{\sqrt{M}} \sum_{m=1}^{M} \epsilon_m \nabla f_m\left(\boldsymbol{x}_\star\right) \right\|_2^2} \\
&= \sqrt{\frac{1}{M} \mathbb{E} \sum_{m=1}^{M} \|\nabla f_m\left(\boldsymbol{x}_\star\right)\|_2^2} \\
&= \sqrt{\mathbb{E}\,\|\nabla f\left(\boldsymbol{x}_\star\right)\|_2^2},
\end{aligned}$$

where the second and third lines respectively hold because ϵ_ms are independent zero-mean random variables and f_ms are i.i.d. copies of f. The following corollary of Theorem 1 is the immediate result of the bound above.

Corollary 1. *Let $\mathcal{A}_\delta$ and $p_\tau(\mathcal{A}_\delta)$ be defined by (4) and (6), respectively. For any $t>0$, if*

$$M > \left(\frac{2\sqrt{\mathbb{E}\left\|\nabla f(\boldsymbol{x}_\star)\right\|_2^2}+t}{\tau p_\tau(\mathcal{A}_\delta)}\right)^2,$$

then with probability $\geq 1-e^{-2Mt^2/\tau^2}$ any solution $\widehat{\boldsymbol{x}}$ of (3) obeys

$$\left\|\widehat{\boldsymbol{x}}-\boldsymbol{x}_\star\right\|_2 \leq \frac{\frac{1}{M}\sum_{m=1}^M|\xi_m|+\varepsilon}{\tau p_\tau(\mathcal{A}_\delta)-\frac{2\sqrt{\mathbb{E}\|\nabla f(\boldsymbol{x}_\star)\|_2^2}+t}{\sqrt{M}}}.$$

Sparse regression: Next, we approximate $\mathfrak{C}_M(\mathcal{A}_\delta)$ for the case of ℓ_1-regularized anchored regression, i.e., $\Omega(\boldsymbol{x})=\lambda\|\boldsymbol{x}\|_1$ for some $\lambda>0$. Let $\mathcal{S}_\star$ denote the support set of $\boldsymbol{x}_\star$ (i.e., $\mathcal{S}_\star=\{i\in[1,N]:x_{\star i}\neq 0\}$) and $s=|\mathcal{S}_\star|$. The subdifferential of the $\Omega(\cdot)$ at $\boldsymbol{x}_\star$ can be expressed as

$$\partial\Omega(\boldsymbol{x}_\star)=\left\{\boldsymbol{g}:\|\boldsymbol{g}\|_\infty\leq\lambda \text{ and } \boldsymbol{g}|_{\mathcal{S}_\star}=\lambda\operatorname{sgn}\left(\boldsymbol{x}_\star|_{\mathcal{S}_\star}\right)\right\}.$$

Therefore, we have

$$\mathcal{A}_\delta\subseteq\left\{\boldsymbol{h}:\sqrt{2(1-\delta)}\|\boldsymbol{h}\|_2-\lambda\left\|\boldsymbol{h}|_{\mathcal{S}_\star^{\mathrm{c}}}\right\|_1+\langle\frac{\boldsymbol{x}_\star|_{\mathcal{S}_\star}}{\|\boldsymbol{x}_\star\|_2}-\lambda\operatorname{sgn}\left(\boldsymbol{x}_\star|_{\mathcal{S}_\star}\right),\boldsymbol{h}|_{\mathcal{S}_\star}\rangle\geq 0\right\}.$$

It follows from the Cauchy-Schwarz inequality that

$$\begin{aligned}\langle\frac{\boldsymbol{x}_\star|_{\mathcal{S}_\star}}{\|\boldsymbol{x}_\star\|_2}-\lambda\operatorname{sgn}\left(\boldsymbol{x}_\star|_{\mathcal{S}_\star}\right),\boldsymbol{h}|_{\mathcal{S}_\star}\rangle &\leq \left\|\frac{\boldsymbol{x}_\star|_{\mathcal{S}_\star}}{\|\boldsymbol{x}_\star\|_2}-\lambda\operatorname{sgn}\left(\boldsymbol{x}_\star|_{\mathcal{S}_\star}\right)\right\|_2\left\|\boldsymbol{h}|_{\mathcal{S}_\star}\right\|_2\\ &=\left\|\frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2}-\lambda\operatorname{sgn}(\boldsymbol{x}_\star)\right\|_2\left\|\boldsymbol{h}|_{\mathcal{S}_\star}\right\|_2,\end{aligned}$$

thereby

$$\mathcal{A}_\delta\subseteq\mathcal{A}_\delta'\overset{\text{def}}{=}\left\{\boldsymbol{h}:\sqrt{2(1-\delta)}\|\boldsymbol{h}\|_2-\lambda\left\|\boldsymbol{h}|_{\mathcal{S}_\star^{\mathrm{c}}}\right\|_1+\left\|\frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2}-\lambda\operatorname{sgn}(\boldsymbol{x}_\star)\right\|_2\left\|\boldsymbol{h}|_{\mathcal{S}_\star}\right\|_2\geq 0\right\}.$$

For any non-zero $\boldsymbol{h}$ and any $\boldsymbol{z}\in\mathbb{R}^N$, it follows from Hölder's inequality that

$$\langle\boldsymbol{z},\frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\leq\left\|\boldsymbol{z}|_{\mathcal{S}_\star}\right\|_2\frac{\left\|\boldsymbol{h}|_{\mathcal{S}_\star}\right\|_2}{\|\boldsymbol{h}\|_2}+\left\|\boldsymbol{z}|_{\mathcal{S}_\star^{\mathrm{c}}}\right\|_\infty\frac{\left\|\boldsymbol{h}|_{\mathcal{S}_\star^{\mathrm{c}}}\right\|_1}{\|\boldsymbol{h}\|_2}. \tag{8}$$

Therefore, recalling the definition (5), we can write

$$\begin{aligned}\mathfrak{C}_M(\mathcal{A}_\delta)&\leq\mathfrak{C}_M(\mathcal{A}_\delta')\\ &=\mathbb{E}\sup_{\boldsymbol{h}\in\mathcal{A}_\delta'}\frac{1}{\sqrt{M}}\sum_{m=1}^M\epsilon_m\langle\nabla f_m(\boldsymbol{x}_\star),\frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\\ &\leq\mathbb{E}\sup_{\boldsymbol{h}\in\mathcal{A}_\delta'}\left\|\frac{1}{\sqrt{M}}\sum_{m=1}^M\epsilon_m\nabla f_m(\boldsymbol{x}_\star)|_{\mathcal{S}_\star}\right\|_2\frac{\left\|\boldsymbol{h}|_{\mathcal{S}_\star}\right\|_2}{\|\boldsymbol{h}\|_2}\\ &\quad+\mathbb{E}\sup_{\boldsymbol{h}\in\mathcal{A}_\delta'}\left\|\frac{1}{\sqrt{M}}\sum_{m=1}^M\epsilon_m\nabla f_m(\boldsymbol{x}_\star)|_{\mathcal{S}_\star^{\mathrm{c}}}\right\|_\infty\frac{\left\|\boldsymbol{h}|_{\mathcal{S}_\star^{\mathrm{c}}}\right\|_1}{\|\boldsymbol{h}\|_2}\end{aligned}$$

where the first inequality follows from the fact that $\mathcal{A}_\delta \subseteq \mathcal{A}'_\delta$ and the second inequality is an application of (8) for $\boldsymbol{z} = \frac{1}{\sqrt{M}} \sum_{m=1}^{M} \epsilon_m \nabla f_m(\boldsymbol{x}_\star)$ followed by a triangle inequality. We always have $\frac{\left\|\boldsymbol{h}|_{\mathcal{S}_\star}\right\|_2}{\|\boldsymbol{h}\|_2} \leq 1$ and the definition of $\mathcal{A}'_\delta$ implies that

$$\begin{aligned}\frac{\left\|\boldsymbol{h}|_{\mathcal{S}_\star^{\mathrm{c}}}\right\|_1}{\|\boldsymbol{h}\|_2} &\leq \frac{\sqrt{2(1-\delta)}}{\lambda} + \frac{1}{\lambda}\left\|\frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2} - \lambda \operatorname{sgn}(\boldsymbol{x}_\star)\right\|_2 \frac{\left\|\boldsymbol{h}|_{\mathcal{S}_\star}\right\|_2}{\|\boldsymbol{h}\|_2} \\ &\leq \frac{\sqrt{2(1-\delta)}}{\lambda} + \frac{1}{\lambda}\left\|\frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2} - \lambda \operatorname{sgn}(\boldsymbol{x}_\star)\right\|_2 .\end{aligned}$$

Hence, we deduce that

$$\begin{aligned}\mathfrak{C}_M(\mathcal{A}_\delta) \leq\ & \mathbb{E}\left\|\frac{1}{\sqrt{M}}\sum_{m=1}^{M} \epsilon_m \nabla f_m(\boldsymbol{x}_\star)|_{\mathcal{S}_\star}\right\|_2 \\ &+ \left(\frac{\sqrt{2(1-\delta)}}{\lambda} + \frac{1}{\lambda}\left\|\frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2} - \lambda \operatorname{sgn}(\boldsymbol{x}_\star)\right\|_2\right) \mathbb{E}\left\|\frac{1}{\sqrt{M}}\sum_{m=1}^{M} \epsilon_m \nabla f_m(\boldsymbol{x}_\star)|_{\mathcal{S}_\star^{\mathrm{c}}}\right\|_\infty .\end{aligned}$$

Next, we bound the terms with the expectation. For the first term, the Cauchy-Schwarz inequality and the fact that ϵ_ms are independent and zero-mean guarantees that

$$\begin{aligned}\mathbb{E}\left\|\frac{1}{\sqrt{M}}\sum_{m=1}^{M} \epsilon_m \nabla f_m(\boldsymbol{x}_\star)|_{\mathcal{S}_\star}\right\|_2 &\leq \sqrt{\mathbb{E}\left\|\frac{1}{\sqrt{M}}\sum_{m=1}^{M} \epsilon_m \nabla f_m(\boldsymbol{x}_\star)|_{\mathcal{S}_\star}\right\|_2^2} \\ &= \sqrt{\frac{1}{M}\sum_{m=1}^{M} \mathbb{E}\left\|\nabla f_m(\boldsymbol{x}_\star)|_{\mathcal{S}_\star}\right\|_2^2} \\ &= \sqrt{\mathbb{E}\left\|\nabla f(\boldsymbol{x}_\star)|_{\mathcal{S}_\star}\right\|_2^2}.\end{aligned}$$

For the second term, a similar application of the Cauchy-Schwarz inequality, followed by the Nemirovsky's inequality (see [18, Lemma 5.2.2] and [2, Theorem 2.2]) with the minor assumption that $N - s \geq 3$, yield

$$\begin{aligned}\mathbb{E}\left\|\frac{1}{\sqrt{M}}\sum_{m=1}^{M} \epsilon_m \nabla f_m(\boldsymbol{x}_\star)|_{\mathcal{S}_\star^{\mathrm{c}}}\right\|_\infty &\leq \sqrt{\mathbb{E}\left\|\frac{1}{\sqrt{M}}\sum_{m=1}^{M} \epsilon_m \nabla f_m(\boldsymbol{x}_\star)|_{\mathcal{S}_\star^{\mathrm{c}}}\right\|_\infty^2} \\ &\leq \sqrt{2e\log(N-s)\cdot\frac{1}{M}\sum_{m=1}^{M} \mathbb{E}\left\|\nabla f_m(\boldsymbol{x}_\star)|_{\mathcal{S}_\star^{\mathrm{c}}}\right\|_\infty^2} \\ &\leq \sqrt{2e\log N\, \mathbb{E}\left\|\nabla f(\boldsymbol{x}_\star)|_{\mathcal{S}_\star^{\mathrm{c}}}\right\|_\infty^2}.\end{aligned}$$

Therefore, we obtain the following corollary of Theorem 1.

Corollary 2. *Let $\mathcal{A}_\delta$ and $p_\tau(\mathcal{A}_\delta)$ be defined by (4) and (6), respectively. Furthermore, let $C^\star_{\delta,\lambda} = 2\sqrt{2e}\left(\frac{\sqrt{2(1-\delta)}}{\lambda} + \frac{1}{\lambda}\left\|\frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2} - \lambda \operatorname{sgn}(\boldsymbol{x}_\star)\right\|_2\right)$. For any $t > 0$, if*

$$M > \left(\frac{2\sqrt{\mathbb{E}\left\|\nabla f(\boldsymbol{x}_\star)|_{\mathcal{S}_\star}\right\|_2^2} + C^\star_{\delta,\lambda}\sqrt{\log(N)\,\mathbb{E}\left\|\nabla f(\boldsymbol{x}_\star)|_{\mathcal{S}_\star^c}\right\|_\infty^2} + t}{\tau p_\tau(\mathcal{A}_\delta)}\right)^2,$$

then with probability $\geq 1 - e^{-2Mt^2/\tau^2}$ any solution $\widehat{\boldsymbol{x}}$ of (3) obeys

$$\|\widehat{\boldsymbol{x}} - \boldsymbol{x}_\star\|_2 \leq \frac{\frac{1}{M}\sum_{m=1}^M |\xi_m| + \varepsilon}{\tau p_\tau(\mathcal{A}_\delta) - \frac{2\sqrt{\mathbb{E}\|\nabla f(\boldsymbol{x}_\star)\|_2^2} + C^\star_{\delta,\lambda}\sqrt{\log(N)\,\mathbb{E}\left\|\nabla f(\boldsymbol{x}_\star)|_{\mathcal{S}_\star^c}\right\|_\infty^2} + t}{\sqrt{M}}}.$$

Let us consider the example of linear regression for which we can make the result of Corollary 2 more explicit. Specifically, suppose that $f(\boldsymbol{x}) = \boldsymbol{a}^{\mathsf{T}}\boldsymbol{x}$ for $\boldsymbol{a} \sim \text{Normal}(\boldsymbol{0}, \boldsymbol{I})$. Then, we have $\nabla f(\boldsymbol{x}_\star) = \boldsymbol{a}$ which implies

$$\mathbb{E}\left\|\nabla f(\boldsymbol{x}_\star)|_{\mathcal{S}_\star}\right\|_2^2 = |\mathcal{S}_\star| = s \qquad \text{and} \qquad \mathbb{E}\left\|\nabla f(\boldsymbol{x}_\star)|_{\mathcal{S}_\star^c}\right\|_\infty^2 \lesssim \log(N).$$

Furthermore, with $\lambda = \frac{1}{\sqrt{s}}$ we can write

$$C^\star_{\delta,\lambda} \leq c_\delta \sqrt{s},$$

where c_δ is a constant decreasing in δ. Therefore, it suffices to have

$$M \gtrsim \left(\frac{\sqrt{s} + c_\delta\sqrt{s}\log(N)}{\tau p_\tau(\mathcal{A}_\delta)}\right)^2$$

to meet the requirement in Corollary 2. This result suggests a sample complexity of $O\left(s\,\log^2 N\right)$ which is optimal up to the dependence on $\log N$. In fact, for this example, the bound of Corollary 2 can be improved using a tighter bound for $\mathbb{E}\left\|\frac{1}{\sqrt{M}}\sum_{m=1}^M \epsilon_m \nabla f_m(\boldsymbol{x}_\star)|_{\mathcal{S}_\star^c}\right\|_\infty$.

3 Recipes for Creating Anchors

While it may be assumed that the anchor vector is provided by an oracle, it is more realistic to have a data-driven method to construct the anchor vector. This requires us to impose new assumptions, albeit implicitly, on the class of functions $\mathcal{F}$ where the random samples $f_1, f_2, \ldots, f_M$ are drawn from. A natural assumption is that for a certain sample loss function $\ell : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ the corresponding *risk* $R(\boldsymbol{x}) \stackrel{\text{def}}{=} \mathbb{E}\ell(f(\boldsymbol{x}), f(\boldsymbol{x}_\star))$ "encodes" the information about $\boldsymbol{x}_\star$ in its derivatives at the origin, i.e., $\nabla^k R(\boldsymbol{0})$ for $k = 1, 2, 3, \ldots$. This assumption is reasonable, for instance, if $R(\boldsymbol{x})$ has a Taylor series approximation around the origin. Note that we are implicitly assuming that the random functions drawn from the class $\mathcal{F}$ as well as the sample loss $\ell(\cdot,\cdot)$ are k-times differentiable. Furthermore, to allow the derivative operator and expectation to commute, the required regularity conditions for the law of f is assumed.

Definition 1 (Spiked Derivatives)**.** We say that the risk $R(\boldsymbol{x}) = \mathbb{E}\ell(f(\boldsymbol{x}), f(\boldsymbol{x}_\star))$ has a *spiked* kth*-order derivative* if $\boldsymbol{x}_\star$ is a *simple principal eigenvector* of $-\nabla^k R(\mathbf{0}) \neq \mathbf{0}$. Namely,

$$\frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2} \in \operatorname*{argmax}_{\boldsymbol{u}\in\mathbb{S}^{N-1}} \left\langle -\nabla^k R(\mathbf{0}), \boldsymbol{u}^{\otimes k} \right\rangle, \tag{9}$$

and every other maximizer is parallel to $\boldsymbol{x}_\star$.

For computational considerations, we focus only on risks with *spiked gradient* (i.e., 1st-order derivative) or *spiked Hessian* (i.e., 2nd-order derivative). In the case of spiked gradients the condition (9) reduces to $-\nabla R(\mathbf{0})$ being perfectly aligned with $\boldsymbol{x}_\star$, that is,

$$\frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2} = -\frac{\nabla R(\mathbf{0})}{\|\nabla R(\mathbf{0})\|_2}. \tag{10}$$

Similarly, for spiked Hessians, the condition (9) reduces to $\boldsymbol{x}_\star$ being a *simple principal eigenvector* of $-\nabla^2 R(\mathbf{0})$. Specifically, we have

$$\begin{gathered}\frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2} \in \operatorname*{argmax}_{\boldsymbol{u}\in\mathbb{S}^{N-1}} -\boldsymbol{u}^\mathsf{T}\nabla^2 R(\mathbf{0})\,\boldsymbol{u} \\ \gamma_\star \overset{\text{def}}{=} \lambda_1\left(-\nabla^2 R(\mathbf{0})\right) - \lambda_2\left(-\nabla^2 R(\mathbf{0})\right) > 0\end{gathered}, \tag{11}$$

where $\lambda_i(\cdot)$ denotes the ith largest (multiplicity inclusive) eigenvalue of its argument. For example, with the squared error $\ell(s,t) = \frac{1}{2}(s-t)^2$ as the sample loss function, if the risk $R(\boldsymbol{x}) = \mathbb{E}\,\frac{1}{2}(f(\boldsymbol{x}) - f(\boldsymbol{x}_\star))^2$ has a spiked gradient then

$$\frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2} = \frac{1}{\|\mathbb{E}\left(\left(f(\boldsymbol{x}_\star) - f(\mathbf{0})\right)\nabla f(\mathbf{0})\right)\|_2}\mathbb{E}\left(\left(f(\boldsymbol{x}_\star) - f(\mathbf{0})\right)\nabla f(\mathbf{0})\right).$$

and if it has spiked Hessian then

$$\frac{\boldsymbol{x}_\star}{\|\boldsymbol{x}_\star\|_2} \in \operatorname*{argmax}_{\boldsymbol{u}\in\mathbb{S}^{N-1}} \boldsymbol{u}^\mathsf{T}\mathbb{E}\left(f(\boldsymbol{x}_\star)\nabla^2 f(\mathbf{0}) - \nabla f(\mathbf{0})\,\nabla^\mathsf{T} f(\mathbf{0})\right)\boldsymbol{u},$$

and the Hessian has a positive spectral gap.

Of course, in practice we do not have access to $R(\boldsymbol{x})$. With finite number of observations, however, the empirical risk

$$R_M(\boldsymbol{x}) = \frac{1}{M}\sum_{m=1}^{M} \ell(f_m(\boldsymbol{x}), y_m),$$

may provide a sufficiently good approximation for $R(\boldsymbol{x})$. Therefore, with sufficient number of observations, $\nabla R_M(\mathbf{0})$, as an approximation to $\nabla R(\mathbf{0})$, or $\nabla^2 R_M(\mathbf{0})$, as an approximation to $\nabla^2 R(\mathbf{0})$, can be used to find an anchor vector under the spiked gradient or spiked Hessian conditions, respectively. In case of the spiked gradient, it suffices to have sufficiently large number of samples (i.e., M) such that

$$\boldsymbol{a}_0 = -\frac{\nabla R_M(\mathbf{0})}{\|\nabla R_M(\mathbf{0})\|_2},$$

obeys anchor vectors' required property (2). In the case that $R(\cdot)$ has a spiked Hessian, using a variant of the Davis–Kahan's theorem [24, Corollary 3], we can show that for

$$\boldsymbol{a}_0 \in \underset{\boldsymbol{u}\in\mathbb{S}^{N-1}}{\operatorname{argmax}} \; -\boldsymbol{u}^\top \nabla^2 R_M(\boldsymbol{0})\, \boldsymbol{u}\,, \tag{12}$$

we have

$$\left\| \boldsymbol{a}_0\boldsymbol{a}_0^\top - \frac{1}{\|\boldsymbol{x}_\star\|_2^2}\boldsymbol{x}_\star\boldsymbol{x}_\star^\top \right\| \le \frac{2\,\|\nabla^2 R_M(\boldsymbol{0}) - \nabla^2 R(\boldsymbol{0})\|}{\gamma_\star}\,,$$

where $\gamma_\star$ is the spectral gap defined in (11). Therefore, if $\|\nabla^2 R_M(\boldsymbol{0}) - \nabla^2 R(\boldsymbol{0})\|$ is sufficiently small relative to the spectral gap $\gamma_\star$, the inequality above implies $\boldsymbol{a}_0$ (or $-\boldsymbol{a}_0$) can obey the required property (2) for some $\delta > 0$. Depending on the law of f, we can bound $\|\nabla^2 R_M(\boldsymbol{0}) - \nabla^2 R(\boldsymbol{0})\|$ using matrix concentration inequalities as $\nabla^2 R_M(\boldsymbol{0})$ can be written as a sum of independent random matrices.

How does construction of the anchor change if we need to estimate a structured ground truth (e.g., a sparse vector)? In this scenario, the situation is more complicated compared to what described above and often there is no definitive answer. The ideal goal is to construct the anchor without increasing the sample complexity of the estimator. This may be possible if we impose an explicit structural constraint $\boldsymbol{u} \in \mathcal{S}$ with $\mathcal{S}$ denoting a prescribed set of structured vectors; namely we would have

$$\boldsymbol{a}_0 \in \underset{\boldsymbol{u}\in\mathbb{S}^{N-1},\boldsymbol{u}\in\mathcal{S}}{\operatorname{argmax}} \; \left\langle -\nabla^k R_M(\boldsymbol{0})\,, \boldsymbol{u}^{\otimes k}\right\rangle. \tag{13}$$

However, this anchor, regardless of whether it obeys (2) at an ideal sample complexity, is often hard to compute. Some relaxations of (13) may produce the desired anchor, perhaps at the cost of increasing M well beyond the required sample complexity (7) of the main estimation procedure. As a concrete example, consider phase retrieval for a (real-valued) sparse target $\boldsymbol{x}_\star$ using noiseless measurements obtained through i.i.d. copies of $\boldsymbol{a} \sim \text{Normal}(\boldsymbol{0}, \boldsymbol{I})$ as the measurement vectors. As in the ordinary phase retrieval, the risk $R(\boldsymbol{x}) = \mathbb{E}\frac{1}{2}\left(|\boldsymbol{a}^\top\boldsymbol{x}|^2 - |\boldsymbol{a}^\top\boldsymbol{x}_\star|^2\right)^2$ has a spiked Hessian because $-\nabla^2 R(\boldsymbol{0}) = \mathbb{E}\,|\boldsymbol{a}^\top\boldsymbol{x}_\star|^2\,\boldsymbol{a}\boldsymbol{a}^\top = \|\boldsymbol{x}_\star\|_2^2\,\boldsymbol{I} + 2\boldsymbol{x}_\star\boldsymbol{x}_\star^\top$ whose spectral gap is $\gamma_\star = 2\,\|\boldsymbol{x}_\star\|_2^2$. However, it can be shown that the empirical Hessian $-\nabla^2 R_M(\boldsymbol{0}) = \frac{1}{M}\sum_{m=1}^M |\boldsymbol{a}_m^\top\boldsymbol{x}_\star|^2\,\boldsymbol{a}_m\boldsymbol{a}_m^\top$ does not have a strong concentration around $\nabla^2 R(\boldsymbol{0})$ unless M (i.e., the number of samples) grows quadratically with sparsity of $\boldsymbol{x}_\star$ (i.e., $\|\boldsymbol{x}_\star\|_0$). Therefore, the conventional scheme described in the previous paragraph would lead to a suboptimal sample complexity; ideally M should have linear dependence on $\|\boldsymbol{x}_\star\|_0$. If we have an oracle that can solve *sparse principal component analysis* [25, 9], a computationally intractable procedure, then instead of (12) the sparse principal eigenvector of $-\nabla^2 R_M(\boldsymbol{0})$ can be used as the anchor vector. In fact, this kind of computational oracles effectively solve (13) for $k = 2$ and $\mathcal{S}$ being the set of (sufficiently) sparse vectors and produce the desired $\boldsymbol{a}_0$. This approach would achieve (near) optimal sample complexity as we only need small submatrices of $-\nabla^2 R_M(\boldsymbol{0})$ to concentrate.

4 Proof of the Main Result

In this section we provide a proof of Theorem 1. The argument is based on the small-ball method introduced in [10, 16] with minor modifications. Our derivations mostly follow the

exposition of this method in [22].

Proof of Theorem 1. For $t \geq 0$, let $\psi_t(s) \stackrel{\text{def}}{=} (s)_+ - (s-t)_+$ which is a *contraction* (i.e., $|\psi_\tau(s_2) - \psi_\tau(s_1)| \leq |s_2 - s_1|$ for all $s_1, s_2 \in \mathbb{R}$). We can bound $R_M^+(\boldsymbol{x}_\star + \boldsymbol{h})$ from below as

$$\begin{aligned} R_M^+(\boldsymbol{x}_\star + \boldsymbol{h}) &= \frac{1}{M}\sum_{m=1}^M (f_m(\boldsymbol{x}_\star + \boldsymbol{h}) - y_m)_+ \\ &= \frac{1}{M}\sum_{m=1}^M (f_m(\boldsymbol{x}_\star + \boldsymbol{h}) - f_m(\boldsymbol{x}_\star) - \xi_m)_+ \\ &\geq \frac{1}{M}\sum_{m=1}^M (\langle \nabla f_m(\boldsymbol{x}_\star), \boldsymbol{h}\rangle - \xi_m)_+ \\ &\geq \frac{1}{M}\sum_{m=1}^M (\langle \nabla f_m(\boldsymbol{x}_\star), \boldsymbol{h}\rangle)_+ - \frac{1}{M}\sum_{m=1}^M (\xi_m)_+ \\ &\geq \frac{1}{M}\sum_{m=1}^M \psi_{\tau\|\boldsymbol{h}\|_2}(\langle \nabla f_m(\boldsymbol{x}_\star), \boldsymbol{h}\rangle) - \frac{1}{M}\sum_{m=1}^M (\xi_m)_+ , \end{aligned} \tag{14}$$

where the inequalities hold, respectively, because $f_m(\cdot)$ is convex, $(\cdot)_+$ is subadditive, and $(s)_+ \geq \psi_t(s)$. Furthermore, using the fact that $t\mathbb{1}(s \geq t) \leq \psi_t(s)$, we have

$$\begin{aligned} \tau\|\boldsymbol{h}\|_2 \mathbb{P}(\langle \nabla f(\boldsymbol{x}_\star), \boldsymbol{h}\rangle \geq \tau\|\boldsymbol{h}\|_2) &= \tau\|\boldsymbol{h}\|_2 \mathbb{E}(\mathbb{1}(\langle \nabla f(\boldsymbol{x}_\star), \boldsymbol{h}\rangle \geq \tau\|\boldsymbol{h}\|_2)) \\ &\leq \mathbb{E}\left(\psi_{\tau\|\boldsymbol{h}\|_2}(\langle \nabla f(\boldsymbol{x}_\star), \boldsymbol{h}\rangle)\right) \end{aligned}$$

Therefore, by adding and subtracting the sides of the inequality above in the right-hand side of (14) and some rearrangement we obtain

$$\begin{aligned} R_M^+(\boldsymbol{x}_\star + \boldsymbol{h}) \geq \tau\|\boldsymbol{h}\|_2 \mathbb{P}(\langle \nabla f(\boldsymbol{x}_\star), \boldsymbol{h}\rangle \geq \tau\|\boldsymbol{h}\|_2) - \frac{1}{M}\sum_{m=1}^M (\xi_m)_+ \\ + \frac{1}{M}\sum_{m=1}^M \psi_{\tau\|\boldsymbol{h}\|_2}(\langle \nabla f_m(\boldsymbol{x}_\star), \boldsymbol{h}\rangle) - \mathbb{E}\left(\psi_{\tau\|\boldsymbol{h}\|_2}(\langle \nabla f(\boldsymbol{x}_\star), \boldsymbol{h}\rangle)\right) . \end{aligned} \tag{15}$$

We only need to establish a uniform lower bound over $\boldsymbol{h} \in \mathcal{A}_\delta$ for the expression in the second line of (15). It is easy to verify that for every $\alpha, t \geq 0$ and $s \in \mathbb{R}$ the identity $\psi_{\alpha t}(s) = t\psi_\alpha\left(\frac{s}{t}\right)$ holds.[2] Thus, we write

$$\begin{aligned} &\frac{1}{M}\sum_{m=1}^M \psi_{\tau\|\boldsymbol{h}\|_2}(\langle \nabla f_m(\boldsymbol{x}_\star), \boldsymbol{h}\rangle) - \mathbb{E}\left(\psi_{\tau\|\boldsymbol{h}\|_2}(\langle \nabla f(\boldsymbol{x}_\star), \boldsymbol{h}\rangle)\right) \\ &= -\frac{1}{M}\|\boldsymbol{h}\|_2 \sum_{m=1}^M \mathbb{E}\left(\psi_\tau\left(\langle \nabla f(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\right)\right) - \psi_\tau\left(\langle \nabla f_m(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\right) , \end{aligned}$$

and we only need to find an upper bound for

$$F_{\mathcal{A}_\delta}(f_1, f_2, \ldots, f_M) \stackrel{\text{def}}{=} \sup_{\boldsymbol{h}\in\mathcal{A}_\delta} \frac{1}{M}\sum_{m=1}^M \mathbb{E}\left(\psi_\tau\left(\langle \nabla f(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\right)\right) - \psi_\tau\left(\langle \nabla f_m(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\right) .$$

[2] Because $\psi_\alpha(\cdot)$ is bounded, we can treat $t = 0$ as $t \to 0$ to avoid the issue of division by zero.

Since $\psi_\tau(\cdot)$ is bounded by 0 and τ, a standard application of the *bounded difference inequality* [15] to $F_{\mathcal{A}_\delta}(\cdot)$ shows that for any $t > 0$ we have

$$F_{\mathcal{A}_\delta}(f_1, f_2, \ldots, f_M) \leq \mathbb{E}F_{\mathcal{A}_\delta}(f_1, f_2, \ldots, f_M) + \frac{t}{\sqrt{M}}, \tag{16}$$

with probability $\geq 1 - e^{-2Mt^2/\tau^2}$. It only remains to upper bound $\mathbb{E}F_{\mathcal{A}_\delta}(f_1, f_2, \ldots, f_M)$. Writing the inner expectation in $\mathbb{E}F_{\mathcal{A}_\delta}(f_1, f_2, \ldots, f_M)$ with respect to $\widetilde{f}_m$s that are i.i.d. copies of f independent of everything else and applying Jensen's inequality we obtain

$$\begin{aligned}
\mathbb{E}F_{\mathcal{A}_\delta}(f_1, f_2, \ldots, f_M) &= \mathbb{E}\sup_{\boldsymbol{h}\in\mathcal{A}_\delta} \frac{1}{M}\sum_{m=1}^{M} \mathbb{E}\left(\psi_\tau\left(\langle\nabla f(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\right)\right) - \psi_\tau\left(\langle\nabla f_m(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\right) \\
&\leq \mathbb{E}\sup_{\boldsymbol{h}\in\mathcal{A}_\delta} \frac{1}{M}\sum_{m=1}^{M} \psi_\tau\left(\langle\nabla \widetilde{f}_m(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\right) - \psi_\tau\left(\langle\nabla f_m(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\right).
\end{aligned}$$

The next step is the standard *symmetrization* argument [see e.g., 23, Lemma 2.3.1]. Since f_m and $\widetilde{f}_m$ are i.i.d., multiplying each summand in the right-hand side of the inequality above by a corresponding $\epsilon_m = \pm 1$ does not change the distribution of the random process and thereby the desired expected value. Take $\epsilon_1, \epsilon_2, \ldots, \epsilon_M$ to be i.i.d. Rademacher random variables independent of everything else. Therefore, we obtain

$$\begin{aligned}
\mathbb{E}F_{\mathcal{A}_\delta}(f_1, f_2, \ldots, f_M) &\leq \mathbb{E}\sup_{\boldsymbol{h}\in\mathcal{A}_\delta}\Bigg[\frac{1}{M}\sum_{m=1}^{M}\epsilon_m\psi_\tau\left(\langle\nabla \widetilde{f}_m(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\right) \\
&\qquad -\epsilon_m\psi_\tau\left(\langle\nabla f_m(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\right)\Bigg] \\
&\leq \mathbb{E}\sup_{\boldsymbol{h}\in\mathcal{A}_\delta}\frac{1}{M}\sum_{m=1}^{M}\epsilon_m\psi_\tau\left(\langle\nabla \widetilde{f}_m(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\right) \\
&\qquad + \mathbb{E}\sup_{\boldsymbol{h}\in\mathcal{A}_\delta}\frac{1}{M}\sum_{m=1}^{M} -\epsilon_m\psi_\tau\left(\langle\nabla f_m(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\right) \\
&\leq 2\mathbb{E}\sup_{\boldsymbol{h}\in\mathcal{A}_\delta}\frac{1}{M}\sum_{m=1}^{M}\epsilon_m\psi_\tau\left(\langle\nabla f_m(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\right).
\end{aligned} \tag{17}$$

Because $\psi_\tau(0) = 0$ and $\psi_\tau(\cdot)$ is a contraction , we can invoke the *Rademacher contraction principle* [13, Theorem 4.12] to show that

$$\begin{aligned}
\mathbb{E}\sup_{\boldsymbol{h}\in\mathcal{A}_\delta}\frac{1}{M}\sum_{m=1}^{M}\epsilon_m\psi_\tau\left(\langle\nabla f_m(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle\right) &\leq \mathbb{E}\sup_{\boldsymbol{h}\in\mathcal{A}_\delta}\frac{1}{M}\sum_{m=1}^{M}\epsilon_m\langle\nabla f_m(\boldsymbol{x}_\star), \frac{\boldsymbol{h}}{\|\boldsymbol{h}\|_2}\rangle \\
&= \frac{1}{\sqrt{M}}\mathfrak{C}_M(\mathcal{A}_\delta),
\end{aligned}$$

where $\mathfrak{C}_M(\mathcal{A}_\delta)$ is defined in (5) and can be interpreted as a measure of *complexity* of $\mathcal{A}_\delta$ with respect to the law of f. It follows from (17) that

$$\mathbb{E}F_{\mathcal{A}_\delta}(f_1, f_2, \ldots, f_M) \leq \frac{2}{\sqrt{M}}\mathfrak{C}_M(\mathcal{A}_\delta).$$

This bound together with (16) guarantees that

$$F_{\mathcal{A}_\delta}\left(f_1, f_2, \ldots, f_M\right) \leq \frac{2}{\sqrt{M}}\, \mathfrak{C}_M\left(\mathcal{A}_\delta\right) + \frac{t}{\sqrt{M}}\,,$$

with probability $\geq 1 - e^{-2Mt^2/\tau^2}$. Finally, on the same event, it follows from (15) that for all $\boldsymbol{h} \in \mathcal{A}_\delta$ we have

$$\begin{aligned} R_M^+\left(\boldsymbol{x}_\star + \boldsymbol{h}\right) &\geq \tau \left\|\boldsymbol{h}\right\|_2 \mathbb{P}\left(\left\langle \nabla f\left(\boldsymbol{x}_\star\right), \boldsymbol{h}\right\rangle \geq \tau \left\|\boldsymbol{h}\right\|_2\right) - \frac{1}{M}\sum_{m=1}^M \left(\xi_m\right)_+ \\ &\quad - \frac{2\left\|\boldsymbol{h}\right\|_2}{\sqrt{M}}\, \mathfrak{C}_M\left(\mathcal{A}_\delta\right) - \frac{t\left\|\boldsymbol{h}\right\|_2}{\sqrt{M}} \\ &\geq \tau \left\|\boldsymbol{h}\right\|_2 p_\tau\left(\mathcal{A}_\delta\right) - \frac{1}{M}\sum_{m=1}^M \left(\xi_m\right)_+ - \frac{2\left\|\boldsymbol{h}\right\|_2}{\sqrt{M}}\, \mathfrak{C}_M\left(\mathcal{A}_\delta\right) - \frac{t\left\|\boldsymbol{h}\right\|_2}{\sqrt{M}}\,. \end{aligned}$$

If $\boldsymbol{x}_\star + \boldsymbol{h}$ is feasible in (3), then

$$\begin{aligned} R_M^+\left(\boldsymbol{x}_\star + \boldsymbol{h}\right) &\leq R_M^+\left(\boldsymbol{x}_\star\right) + \varepsilon \\ &= \frac{1}{M}\sum_{m=1}^M \left(-\xi_m\right)_+ + \varepsilon\,. \end{aligned}$$

Therefore, if $\widehat{\boldsymbol{x}}$ denotes a solution to (3) and $\boldsymbol{h} = \widehat{\boldsymbol{x}} - \boldsymbol{x}_\star$, then we have

$$\begin{aligned} &\tau \left\|\boldsymbol{h}\right\|_2 p_\tau\left(\mathcal{A}_\delta\right) - \frac{1}{M}\sum_{m=1}^M \left(\xi_m\right)_+ \\ &\quad - \frac{2\left\|\boldsymbol{h}\right\|_2}{\sqrt{M}}\, \mathfrak{C}_M\left(\mathcal{A}_\delta\right) - \frac{t\left\|\boldsymbol{h}\right\|_2}{\sqrt{M}} \leq \frac{1}{M}\sum_{m=1}^M \left(-\xi_m\right)_+ + \varepsilon\,, \end{aligned}$$

or equivalently

$$\left(\tau p_\tau\left(\mathcal{A}_\delta\right) - \frac{2\mathfrak{C}_M\left(\mathcal{A}_\delta\right) + t}{\sqrt{M}}\right)\left\|\boldsymbol{h}\right\|_2 \leq \frac{1}{M}\sum_{m=1}^M \left|\xi_m\right| + \varepsilon\,.$$

The proof is complete because of the assumption $M > \left(\frac{2\mathfrak{C}_M(\mathcal{A}_\delta)+t}{\tau p_\tau(\mathcal{A}_\delta)}\right)^2$. □

References

[1] S. Bahmani and J. Romberg. Phase retrieval meets statistical learning theory: A flexible convex relaxation. to appear at AISTATS'17, preprint arXiv: 1610.04210 [cs.IT], Oct. 2016.

[2] L. Dümbgen, S. A. van de Geer, M. C. Veraar, and J. A. Wellner. Nemirovski's inequalities revisited. *American mathematical monthly*, 117(2):138–160, 2010.

[3] T. Goldstein and C. Studer. PhaseMax: convex phase retrieval via basis pursuit. preprint arXiv: 1610.07531 [cs.IT], Oct. 2016.

[4] P. Hand and V. Voroninski. Compressed sensing from phaseless gaussian measurements via linear programming in the natural parameter space. preprint arXiv:1611.05985, 2016.

[5] P. Hand and V. Voroninski. Corruption robust phase retrieval via linear programming. preprint arXiv:1612.03547, 2016.

[6] P. Hand and V. Voroninski. An elementary proof of convex phase retrieval in the natural parameter space via the linear program PhaseMax. preprint arXiv:1611.03935, 2016.

[7] S. Haykin. *Neural Networks and Learning Machines*. Pearson, Upper Saddle River, NJ, USA, 3rd edition, 2009.

[8] H. Ichimura. Semiparametric least squares (SLS) and weighted SLS estimation of single-index models. *Journal of Econometrics*, 58(1):71 – 120, 1993.

[9] I. M. Johnstone and A. Y. Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486): 682–693, 2009.

[10] V. Koltchinskii and S. Mendelson. Bounding the smallest singular value of a random matrix without concentration. *International Mathematics Research Notices*, 2015(23): 12991–13008, 2015.

[11] G. Lecué and S. Mendelson. Regularization and the small-ball method I: Sparse recovery. preprint arXiv:1601.05584 [math.ST], 2016.

[12] G. Lecué and S. Mendelson. Regularization and the small-ball method II: Complexity dependent error rates. preprint arXiv:1608.07681 [math.ST], 2016.

[13] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: Isoperimetry and processes*. Springer Science & Business Media, 2013.

[14] P. Mccullagh and J. A. Nelder. *Generalized linear models*, volume 37 of *Monographs on statistics and applied probability*. Chapman and Hall/CRC, London ; New York, 2nd edition, 1989.

[15] C. McDiarmid. On the method of bounded differences. *Surveys in combinatorics*, 141 (1):148–188, 1989.

[16] S. Mendelson. Learning without concentration. In *Proceedings of the 27th Conference on Learning Theory (COLT)*, volume 35 of *JMLR W&CP*, pages 25–39, 2014.

[17] S. Mendelson. Learning without concentration for general loss functions. preprint; arXiv:1410.3192, 2014.

[18] A. Nemirovski. *Topics in Non-parametric Statistics*, chapter 5, pages 183–206. Springer Berlin Heidelberg, Berlin, Heidelberg, 2000.

[19] Y. Plan and R. Vershynin. The generalized LASSO with non-linear observations. *IEEE Transactions on Information Theory*, 62(3):1528–1537, Mar. 2016.

[20] Y. Plan, R. Vershynin, and E. Yudovina. High-dimensional estimation with geometric constraints. *Information and Inference*, 2016.

[21] S. Shalev-Shwartz and S. Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, New York, NY, USA, 2014.

[22] J. A. Tropp. *Convex Recovery of a Structured Signal from Independent Random Linear Measurements*, pages 67–101. Springer International Publishing, Cham, 2015.

[23] A. W. van Der Vaart and J. A. Wellner. *Weak Convergence and Empirical Processes*. Springer Series in Statistics. Springer, 1996.

[24] Y. Yu, T. Wang, and R. J. Samworth. A useful variant of the Davis–Kahan theorem for statisticians. *Biometrika*, 102(2):315–323, 2014.

[25] H. Zou, T. Hastie, and R. Tibshirani. Sparse principal component analysis. *Journal of Computational and Graphical Statistics*, 15(2):265–286, 2006.